

No Agent Left Behind: Dynamic Fair Division of Multiple Resources

Ian Kash
Microsoft Research
Cambridge, UK
iankash@microsoft.com

Ariel D. Procaccia
Carnegie Mellon University,
USA
arielpro@cs.cmu.edu

Nisarg Shah
Carnegie Mellon University,
USA
nkshah@cs.cmu.edu

ABSTRACT

Recently fair division theory has emerged as a promising approach for the allocation of multiple computational resources among agents. While in reality agents are not all present in the system simultaneously, previous work has studied *static* settings where all relevant information is known upfront. Our goal is to better understand the *dynamic* setting. On the conceptual level, we develop a dynamic model of fair division, and propose desirable axiomatic properties for dynamic resource allocation mechanisms. On the technical level, we construct two novel mechanisms that provably satisfy some of these properties, and analyze their performance using real data. We believe that our work informs the design of superior multiagent systems, and at the same time expands the scope of fair division theory by initiating the study of dynamic and fair resource allocation mechanisms.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems;
J.4 [Computer Applications]: Social and Behavioral Sciences—
Economics

General Terms

Algorithms, Economics

Keywords

Resource allocation, Fair division, Mechanism design

1. INTRODUCTION

The question of how to *fairly* divide goods or resources has been the subject of intellectual curiosity for millennia. While early solutions can be traced back to ancient writings, rigorous approaches to fairness were proposed only as late as the mid Twentieth Century, by mathematicians and social scientists. Over time, *fair division* has emerged as an influential subfield of microeconomic theory. In the last few years fair division has also attracted the attention of AI researchers (see, e.g., [4, 12, 2]), who envision applications of fair division in multiagent systems [3]. However, fair division theory has seen few applications to date.

It is only very recently that an exciting combination of technological advances and theoretical innovations has pointed the way

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

towards *concrete* applications of fair division. In modern data centers, clusters, and grids, multiple computational resources (such as CPU, memory, and network bandwidth) must be allocated among heterogeneous agents. Agents' demands for resources is typically highly structured, as we explain below. Several recent papers [9, 8, 11, 6] suggest that classic fair division mechanisms possess excellent properties in these environments, in terms of their fairness guarantees as well as their game-theoretic properties.

Nevertheless, some aspects of realistic computing systems are beyond the current scope of fair division theory. Perhaps most importantly, the literature does not capture the *dynamics* of these systems. Indeed, it is typically not the case that all the agents are present in the system at any given time; agents may arrive and depart, and the system must be able to adjust the allocation of resources. Even on the conceptual level, dynamic settings challenge some of the premises of fair division theory. For example, if one agent arrives before another, the first agent should intuitively have priority; what does fairness mean in this context? We introduce the concepts that are necessary to answer this question, and design novel mechanisms that satisfy our proposed desiderata. Our contribution is therefore twofold: we *design more realistic resource allocation mechanisms for multiagent systems* that provide theoretical guarantees, and at the same time we *expand the scope of fair division theory* to capture dynamic settings.

Overview of model and results

As in previous papers (e.g., [8, 11]), we assume that agents demand the resources in fixed proportions. Such *Leontief preferences*—as they are known in economics—are easily justified in typical settings where agents must run many instances of a single task (e.g., map jobs in the MapReduce framework). Hence, for example, an agent that requires twice as much CPU as RAM to run a task prefers to be allocated 4 CPU units and 2 RAM units to 2 CPU units and 1 RAM unit, but is indifferent between the former allocation and 5 CPU units and 2 RAM units.

We consider environments where agents arrive over time (but do not depart—see Section 7 for additional discussion of this point). We aim to design resource allocation mechanisms that make *irrevocable* allocations, i.e., the mechanism can allocate more resources to an agent over time, but can never take resources back.

We adapt prominent notions of fairness, efficiency, and truthfulness to our dynamic settings. For fairness, we ask for *envy freeness (EF)*, in the sense that agents like their own allocation best; and *sharing incentives (SI)*, so that agents prefer their allocation to their proportional share of the resources. We also seek *strategyproof (SP)* mechanisms: agents cannot gain from misreporting their demands. Finally, we introduce the notion of *dynamic Pareto optimality (DPO)*: if k agents are entitled to k/n of each resource,

the allocation should not be dominated (in a sense that will be formalized later) by allocations that divide these entitlements. Our first result (in Section 3) is an impossibility: DPO and EF are incompatible. We proceed by relaxing each of these properties.

In Section 4, we relax the EF property. The new dynamic property, which we call *dynamic EF (DEF)*, allows an agent to envy another agent that arrived earlier, as long as the former agent was not allocated resources after the latter agent’s arrival. We construct a new mechanism, DYNAMIC DRF, and prove that it satisfies SI, DEF, SP, and DPO.

In Section 5, we relax the DPO property. Our *cautious DPO (CDPO)* notion allows allocations to only compete with allocations that can ultimately guarantee EF, regardless of the demands of future agents. We design a mechanism called CAUTIOUS LP, and show that it satisfies SI, EF, SP, and CDPO. In a sense, our theoretical results are tight: EF and DPO are incompatible, but relaxing only one of these two properties is sufficient to enable mechanisms that satisfy both, in conjunction with SI and SP.

Despite the assumptions imposed by our theoretical model, we believe that our new mechanisms are compelling, useful guides for the design of practical resource allocation mechanisms in realistic settings. Indeed, in Section 6, we test our mechanisms on real data obtained from a trace of workloads on a Google cluster, and obtain encouraging results.

Related work

Walsh [13] proposed the problem of fair cake cutting where agents arrive, take a piece of cake, and immediately depart. The cake cutting setting deals with the allocation of a *single, heterogeneous* divisible resource; contrast with our setting, which deals with *multiple, homogeneous* divisible resources. Walsh suggested several desirable properties for cake cutting mechanisms in this setting, and showed that adaptations of classic mechanisms achieve these properties (Walsh also pointed out that allocating the whole cake to the first agent achieves the same properties). In particular, his notion of *forward envy freeness*, which is discussed below, is related to our notion of *dynamic envy freeness*.

The networking community has studied the problem of fairly allocating a single homogeneous resource in a queuing model where each agent’s task requires a given number of time units to be processed. In other words, in these models tasks are processed over time, but demands stay fixed, and there are no other dynamics such as agent arrivals and departures. The well-known *fair queuing* [5] solution allocates one unit per agent in successive round-robin fashion. This solution has also been analyzed by economists [10].

Previous papers on the allocation of multiple resources study a static setting. For example, Ghodsi et al. [8] proposed the *dominant resource fairness (DRF)* mechanism, which guarantees a number of desirable theoretical properties. Gutman and Nisan [9] considered generalizations of DRF in a more general model of utilities, and also gave a polynomial time algorithm for another mechanism that was constructed by Dolev et al. [6]. Parkes et al. [11] extended DRF in several ways, and in particular studied the case of indivisible tasks. Most recently, Ghodsi et al. [7] extended DRF to the queuing domain. We elaborate on several of these results below.

2. PRELIMINARIES

In our setting, each agent has a task that requires fixed amounts of different resources. The utility of the agent depends on the quantity (possibly fractional) of its tasks that it can execute given the allocated resources. Formally, denote the set of agents by $N = \{1, \dots, n\}$, and the set of resources by R , $|R| = m$. Let D_{ir} denote the ratio between the maximum amount of resource r agent i

can use given the amounts of other resources present in the system and the total amount of that resource available in the system, either allocated or free. In other words, D_{ir} is the *fraction* of resource r required by agent i . Following [8], the *dominant resource* of agent i is defined as the resource r that maximizes D_{ir} , and the fraction of dominant resource allocated to agent i is called its *dominant share*. Following [11], the (normalized) *demand vector* of agent i is given by $\mathbf{d}_i = \langle d_{i1}, \dots, d_{im} \rangle$, where $d_{ir} = D_{ir} / (\max_{r'} D_{ir'})$ for each resource r . Let \mathcal{D} be the set of all possible normalized demand vectors. Let $\mathbf{d}_{\leq k} = \langle \mathbf{d}_1, \dots, \mathbf{d}_k \rangle$ denote the demand vectors of agents 1 through k . Similarly, let $\mathbf{d}_{>k} = \langle \mathbf{d}_{k+1}, \dots, \mathbf{d}_n \rangle$ denote the demand vectors of agents $k + 1$ through n .

An allocation \mathbf{A} allocates a fraction A_{ir} of resource r to agent i , subject to the feasibility condition $\sum_{i \in N} A_{ir} \leq 1$ for all $r \in R$. Throughout the paper we assume that resources are divisible and that each agent requires a positive amount of each resource, i.e., $d_{ir} > 0$ for all $i \in N$ and $r \in R$. Under such allocations, our model for preferences coincides with the domain of *Leontief preferences*, where the utility of an agent for its allocation vector \mathbf{A}_i is given by

$$u_i(\mathbf{A}_i) = \max\{y \in \mathbb{R}_+ : \forall r \in R, A_{ir} \geq y \cdot d_{ir}\}.$$

In words, the utility of an agent is the fraction of its dominant resource that it can actually use, given its proportional demands and its allocation of the various resources. However, we *do not* rely on an interpersonal comparison of utilities; an agent’s utility function simply induces ordinal preferences over allocations, and its exact value is irrelevant.

We say that an allocation \mathbf{A} is Pareto dominated by another allocation \mathbf{A}' if $u_i(\mathbf{A}'_i) \geq u_i(\mathbf{A}_i)$ for every agent i , and $u_j(\mathbf{A}'_j) > u_j(\mathbf{A}_j)$ for some agent j . For allocations \mathbf{A} over agents in $S \subseteq N$ and \mathbf{A}' over agents in $T \subseteq N$ such that $S \subseteq T$, we say that \mathbf{A}' is an extension of \mathbf{A} to T if $A'_{ir} \geq A_{ir}$ for every agent $i \in S$ and every resource r . When $S = T$, we simply say that \mathbf{A}' is an extension of \mathbf{A} .

An allocation \mathbf{A} is called *non-wasteful* if for every agent i there exists $y \in \mathbb{R}_+$ such that for all $r \in R$, $A_{ir} = y \cdot d_{ir}$. Note that for a non-wasteful allocation, the utility of an agent is the share of its dominant resource allocated to the agent. Also, if \mathbf{A} is a non-wasteful allocation then for all $i \in N$,

$$u_i(\mathbf{A}'_i) > u_i(\mathbf{A}_i) \Rightarrow \forall r \in R, A'_{ir} > A_{ir}. \quad (1)$$

3. DYNAMIC RESOURCE ALLOCATION: A NEW MODEL

We consider a dynamic resource allocation model where agents arrive at different times and do not depart (see Section 7 for a discussion of this point). We assume that agent 1 arrives first, then agent 2, and in general agent k arrives after agents 1, \dots , $k - 1$; we say that agent k arrives in *step k* . An agent reports its demand when it arrives and the demand does not change over time. Thus, at step k , demand vectors $\mathbf{d}_{\leq k}$ are known, and demand vectors $\mathbf{d}_{>k}$ are unknown. A *dynamic resource allocation mechanism* operates as follows. At each step k , the mechanism takes as input the reported demand vectors $\mathbf{d}_{\leq k}$ and outputs an allocation \mathbf{A}^k over the agents present in the system. Crucially, we assume that allocations are *irrevocable*, i.e., $A_{ir}^k \geq A_{ir}^{k-1}$ for every step $k \geq 2$, every agent $i \leq k - 1$, and every resource r . We assume that the dynamic allocation mechanism knows the total number of agents n upfront. This is the case in real-world systems where agents have contributed resources and these contributions are recorded, e.g., federated clouds and research clusters.

Previous work on static resource allocation (e.g., [8, 11]) focused on designing mechanisms that satisfy four prominent desiderata. Three of these—two fairness properties and one game-theoretic property—immediately extend to the dynamic setting.

1. *Sharing Incentives (SI)*. A dynamic allocation mechanism is SI if $u_i(\mathbf{A}_i^k) \geq u_i(\langle 1/n, \dots, 1/n \rangle)$ for all steps k and all agents $i \leq k$. In words, when an agent arrives it receives an allocation that it likes at least as much as an equal split of the resources. This models a setting where agents have made equal contributions to the system and hence have equal entitlements.
2. *Envy Freeness (EF)*. A dynamic allocation mechanism is EF if $u_i(\mathbf{A}_i^k) \geq u_i(\mathbf{A}_j^k)$ for all steps k and all agents $i, j \leq k$, that is, an agent that is present would never prefer the allocation of another agent.
3. *Strategyproofness (SP)*. A dynamic allocation mechanism is SP if no agent can misreport its demand vector and be strictly better off at any step k , regardless of the reported demands of other agents. Formally, a dynamic allocation mechanism is SP if for any agent $i \in N$ and any step k , if \mathbf{A}_i^k is the allocation to agent i at step k when agent i reports its true demand vector and \mathbf{B}_i^k is the allocation to agent i at step k when agent i reports a different demand vector, then $u_i(\mathbf{A}_i^k) \geq u_i(\mathbf{B}_i^k)$. We avoid introducing additional notations that will not be required later.

In the static setting, the fourth prominent axiom, *Pareto optimality (PO)*, means that the mechanism’s allocation is not Pareto dominated by any other allocation. Of course, in the dynamic setting it is unreasonable to expect the allocation in early stages to be Pareto undominated as we need to save resources for future arrivals (recall that allocations are irrevocable). We believe though that the following definition naturally extends PO to our dynamic setting.

4. *Dynamic Pareto Optimality (DPO)*. A dynamic allocation mechanism is DPO if at each step k , the allocation \mathbf{A}^k returned by the mechanism is not Pareto dominated by any other allocation \mathbf{B}^k that allocates up to a (k/n) -fraction of each resource among the k agents present in the system. Put another way, at each step the allocation should not be Pareto dominated by any other allocation that only redistributes the collective entitlements of the agents present in the system among those agents.

It is straightforward to verify that a non-wasteful mechanism (a mechanism returning a non-wasteful allocation at each step) satisfies DPO if and only if the allocation returned by the mechanism at each step k uses at least a (k/n) -fraction of at least one resource (the assumption of strictly positive demands plays a role here).

Impossibility Result

Ideally, we would like to design a dynamic allocation mechanism that is SI, EF, SP, and DPO. However, we show that even satisfying EF and DPO simultaneously is impossible.

THEOREM 1. *Let $n \geq 3$ and $m \geq 2$. Then no dynamic resource allocation mechanism satisfies EF and DPO.*

PROOF. Consider a setting with three agents and two resources. Agents 1 and 2 have demand vectors $\langle 1, 1/9 \rangle$ and $\langle 1/9, 1 \rangle$, respectively (i.e., $d_{11} = 1$, $d_{12} = 1/9$, etc.). At step 2 (after the second agent arrives), at least one of the two agents must be allocated at

least an $x = 3/5$ share of its dominant resource, otherwise the total fraction each resource that is allocated at step 2 would be less than $x + x \cdot (1/9) = 2/3$, violating DPO. Without loss of generality, assume that agent 1 is allocated at least an $x = 3/5$ share of its dominant resource (resource 1) at step 2. If agent 3 reports the demand vector $\langle 1, 1/9 \rangle$ —identical to that of agent 1—then it can be allocated at most a $2/5$ share of its dominant resource (resource 1), and would envy agent 1.

It is easy to extend this argument to the case of $n > 3$, by adding $n - 3$ agents with demand vectors identical to the demand vector of agent 3. To extend to the case of $m > 2$, let all agents have negligibly small demands for the additional resources. \square

It is interesting to note that if either EF or DPO is dropped, the remaining three axioms can be easily satisfied. For example, the trivial mechanism EQUAL SPLIT that just gives every agent a $1/n$ share of each resource when it arrives satisfies SI, EF and SP. To achieve SI, DPO and SP, consider the following mechanism, which we call DYNAMIC DICTATORSHIP. At each step k , the mechanism allocates a $1/n$ share of each resource to agent k , takes back the shares of different resources that the agent cannot use, and then keeps allocating resources to agent 1 (the dictator) until a k/n share of at least one resource is allocated. Note that the mechanism trivially satisfies SI because it allocates resources as valuable as an equal split to each agent as soon as it arrives. The mechanism satisfies DPO because it is non-wasteful and at every step k it allocates a k/n fraction of at least one resource. It can be easily verified that DYNAMIC DICTATORSHIP is also SP.

Neither EQUAL SPLIT nor DYNAMIC DICTATORSHIP is a compelling mechanism. Since these mechanisms are permitted by dropping EF or DPO entirely, we instead explore relaxations of EF and DPO that rule these mechanisms out and permit more compelling mechanisms. Luckily, our mechanisms turn out to satisfy SI and SP as well.

4. RELAXING ENVY FREENESS

Recall that DPO requires a mechanism to allocate at least a k/n fraction of at least one resource at step k , for every $k \in \{1, \dots, n\}$. Thus the mechanism sometimes needs to allocate a large amount of resources to agents arriving early, potentially making it impossible for the mechanism to prevent the late agents from envying the early agents. In other words, when an agent i enters the system it may envy some agent j that arrived before i did; this envy is inevitable in order to be able to satisfy DPO. However, it would be unfair to agent i if agent j were allocated more resources since agent i arrived while i still envied j . To distill this intuition, we introduce the following dynamic version of EF.

- 2'. *Dynamic Envy Freeness (DEF)*. A dynamic allocation mechanism is DEF if at any step an agent i envies an agent j only if j arrived before i did and j has not been allocated any resources since i arrived. Formally, for every $k \in \{1, \dots, n\}$, if $u_i(\mathbf{A}_j^k) > u_i(\mathbf{A}_i^k)$ then $j < i$ and $\mathbf{A}_j^k = \mathbf{A}_j^{i-1}$.

Walsh [13] studied a dynamic cake cutting setting and proposed *forward EF*, which requires that an agent not envy any agent that arrived later. This notion is weaker than DEF because it does not rule out the case where an agent i envies an agent j that arrived earlier and j received resources since i arrived. In our setting, even the trivial mechanism DYNAMIC DICTATORSHIP (see Section 3) satisfies forward EF, but fails to satisfy our stronger notion of DEF.

We next construct a dynamic resource allocation mechanism—DYNAMIC DRF—that achieves the relaxed fairness notion of DEF,

together with SI, DPO, and SP. The mechanism is given as Algorithm 1. Intuitively, at each step k the mechanism starts from the current allocation among the present agents and keeps allocating resources to agents that have the minimum dominant share *at the same rate*, until a k/n fraction of at least one resource is allocated. Always allocating to agents that have the minimum dominant share ensures that agents are not allocated any resources while they are envied. This *water-filling* mechanism is a dynamic adaptation of the dominant resource fairness (DRF) mechanism proposed by Ghodsi et. al. [8]. See Figure 1 for an example.

ALGORITHM 1: DYNAMIC DRF

Data: Demands \mathbf{d}

Result: Allocation \mathbf{A}^k at each step k

$k \leftarrow 1$;

while $k \leq n$ **do**

$\{x_i^k\}_{i=1}^k \leftarrow$ Solution of the LP in the box below;

$A_{ir}^k \leftarrow x_i^k \cdot d_{ir}, \forall i \leq k$;

$k \leftarrow k + 1$;

end

Maximize M^k
subject to
 $x_i^k \geq M^k, \forall i \leq k$
 $x_i^k \geq x_i^{k-1}, \forall i \leq k-1$
 $\sum_{i=1}^k x_i^k \cdot d_{ir} \leq k/n, \forall r \in R$

THEOREM 2. DYNAMIC DRF satisfies SI, DEF, DPO, and SP, and can be implemented in polynomial time.

PROOF. First we show that DYNAMIC DRF satisfies SI. We need to prove that $x_i^k \geq 1/n$ for all agents $i \leq k$ at every step $k \in \{1, \dots, n\}$. We prove this by induction on k . For the base case $k = 1$, it is easy to see that $x_1^1 = 1/n$ and $M^1 = 1/n$ is a solution of the LP of DYNAMIC DRF and hence the optimal solution satisfies $x_1^1 \geq M^1 \geq 1/n$ (in fact, there is an equality). Assume that this is true at step $k-1$ and let us prove the claim for step k , where $k \in \{2, \dots, n\}$. At step k , one feasible solution of the LP is given by $x_i^k = x_i^{k-1}$ for agents $i \leq k-1$, $x_k^k = 1/n$ and $M^k = 1/n$. To see this, note that it trivially satisfies the first two constraints of the LP, because by the induction hypothesis we have $x_i^{k-1} \geq 1/n$ for $i \leq k-1$. Furthermore, in the proposed feasible solution, for any $r \in R$ we have

$$\sum_{i=1}^k x_i^k \cdot d_{ir} = \sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} + \frac{1}{n} \cdot d_{kr} \leq \frac{k-1}{n} + \frac{1}{n} \leq \frac{k}{n},$$

where the first transition follows from the construction of the feasible solution and the second transition holds because $\{x_i^{k-1}\}_{i=1}^{k-1}$ satisfies the LP of step $k-1$, and in particular the third constraint of the LP. Since a feasible solution achieves $M^k = 1/n$, the optimal solution achieves $M^k \geq 1/n$. Thus in the optimal solution $x_i^k \geq M^k \geq 1/n$ for all $i \leq k$, which is the requirement for SI.

Next we show that DYNAMIC DRF satisfies DPO. Observe that at any step k , the third constraint of the LP must be tight for at least one resource in the optimal solution (otherwise every x_i^k along with M^k can be increased by a sufficiently small quantity, contradicting the optimality of M^k). Thus, at each step k the (non-wasteful) mechanism allocates a k/n fraction of at least one resource, which implies that the mechanism satisfies DPO.

To prove that the mechanism satisfies DEF and SP, we first prove several useful lemmas about the allocations returned by the mechanism. In the proof below, M^k and x_i^k refer to the *optimal* solution

of the LP in step k . Furthermore, we assume that $x_i^k = 0$ for agents $i > k$ (i.e., agents not present in the system are not allocated any resources). We begin with the following lemma, which essentially shows that if an agent is allocated some resources in a step using water filling, then the agent's dominant share after the step will be the minimum among the present agents.

LEMMA 3. At every step $k \in \{1, \dots, n\}$, it holds that $x_i^k = \max(M^k, x_i^{k-1})$ for all agents $i \leq k$.

PROOF. Consider any step $k \in \{1, \dots, n\}$. From the first and the second constraints of the LP it is evident that $x_i^k \geq M^k$ and $x_i^k \geq x_i^{k-1}$ (note that $x_i^{k-1} = 0$), thus $x_i^k \geq \max(M^k, x_i^{k-1})$ for all $i \leq k$. Suppose for contradiction that $x_i^k > \max(M^k, x_i^{k-1})$ for some $i \leq k$. Then x_i^k can be reduced by a sufficiently small $\epsilon > 0$ without violating any constraints. This makes the third constraint of the LP loose by at least $\epsilon \cdot d_{ir}$, for every resource $r \in R$. Consequently, the values of x_j^k for $j \neq i$ and M^k can be increased by a sufficiently small $\delta > 0$ without violating the third constraint of the LP. Finally, ϵ (and correspondingly δ) can be chosen to be small enough so that $x_i^k \geq M^k$ is not violated. It follows that the value of M^k can be increased, contradicting the optimality of M^k . \square (Proof of Lemma 3)

Next we show that at each step k , the dominant shares of agents 1 through k are monotonically non-increasing with their time of arrival. This is intuitive because at every step k , agent k enters with zero dominant share and subsequently we perform water filling, hence monotonicity is preserved.

LEMMA 4. For all agents $i, j \in N$ such that $i < j$, we have $x_i^k \geq x_j^k$ at every step $k \in \{1, \dots, n\}$.

PROOF. Fix any two agents $i, j \in N$ such that $i < j$. We prove the lemma by induction on k . The result trivially holds for $k < j$ since $x_j^k = 0$. Assume that $x_i^{k-1} \geq x_j^{k-1}$ where $k \in \{j, \dots, n\}$. At step k , we have $x_i^k = \max(M^k, x_i^{k-1}) \geq \max(M^k, x_j^{k-1}) = x_j^k$, where the first and the last transition follow from Lemma 3 and the second transition follows from our induction hypothesis. \square (Proof of Lemma 4)

The following lemma shows that if agent j has a greater dominant share than agent i at some step, then j must have arrived before i and j must not have been allocated any resources since i arrived. Observe that this is very close to the requirement of DEF.

LEMMA 5. At any step $k \in \{1, \dots, n\}$, if $x_j^k > x_i^k$ for some agents $i, j \leq k$, then $j < i$ and $x_j^k = x_j^{i-1}$.

PROOF. First, note that $j < i$ trivially follows from Lemma 4. Suppose for contradiction that $x_j^k > x_j^{i-1}$ (it cannot be smaller because allocations are irrevocable). Then there exists a step $t \in \{i, \dots, k\}$ such that $x_j^t > x_j^{t-1}$. Now Lemma 3 implies that $x_j^t = M^t \leq x_i^t$, where the last transition follows because x_i^t satisfies the second constraint of the LP at step t (note that $i \leq t$). However, $x_i^t \geq x_i^t$ due to Lemma 4. Thus, $x_j^t = x_i^t$. Now using Lemma 3, $x_j^{t+1} = \max(M^{t+1}, x_j^t) = \max(M^{t+1}, x_i^t) = x_i^{t+1}$. Extending this argument using a simple induction shows that $x_j^{t'} = x_i^{t'}$ for every step $t' \geq t$, in particular, $x_j^k = x_i^k$, contradicting our assumption. \square (Proof of Lemma 5)

We proceed to show that DYNAMIC DRF satisfies DEF. We need to prove that for any step $k \in \{1, \dots, n\}$ and any agents $i, j \leq k$, if agent i envies agent j in step k (i.e., $u_i(\mathbf{A}_j^k) > u_i(\mathbf{A}_i^k)$), then




Figure 1: Allocations returned by DYNAMIC DRF at various steps for 3 agents with demands $\mathbf{d}_1 = \langle 1, 1/2, 3/4 \rangle$, $\mathbf{d}_2 = \langle 1/2, 1, 3/4 \rangle$, and $\mathbf{d}_3 = \langle 1/2, 1/2, 1 \rangle$. Agent 1 receives a $1/3$ share of its dominant resource at step 1. At step 2, water filling drives the dominant shares of agents 1 and 2 up to $4/9$. At step 3, however, agent 3 can only receive a $1/3$ dominant share and the allocations of agents 1 and 2 remain unchanged.

$j < i$ and $x_j^k = x_j^{i-1}$. First, note that $u_i(\mathbf{A}_j^k) > u_i(\mathbf{A}_i^k)$ trivially implies that $x_j^k > x_i^k$, otherwise for the dominant resource r_i^* of agent i , we would have $A_{ir_i^*}^k = x_i^k \geq x_j^k \geq x_j^k \cdot d_{jr_i^*} = A_{jr_i^*}^k$ and agent i would not envy agent j . Now DEF follows from Lemma 5.

To prove that DYNAMIC DRF is SP, suppose for contradiction that an agent $i \in N$ can report an untruthful demand vector \mathbf{d}_i^t such that the agent is strictly better off in at least one step. Let k be the first such step. Denote by \hat{x}_j^t the dominant share of an agent j at step t with manipulation (for agent i , this is the share of the dominant resource of the untruthful demand vector) and similarly, denote by \hat{M}^t the value of M^t in the optimal solution of the LP of step t with manipulation.

LEMMA 6. $\hat{x}_j^k \geq x_j^k$ for every agent $j \leq k$.

PROOF. For any agent j such that $x_j^k > x_i^k$, we have

$$x_j^k = x_j^{i-1} = \hat{x}_j^{i-1} \leq \hat{x}_j^k.$$

Here, the first transition follows from Lemma 5, the second transition holds because manipulation by agent i does not affect the allocation at step $i-1$, and the third transition follows from the LP. For any agent j with $x_j^k \leq x_i^k$, we have

$$x_j^k \leq x_i^k < \hat{x}_i^k = \hat{M}^k \leq \hat{x}_j^k.$$

The second transition is true because if $\hat{x}_i^k \leq x_i^k$ then agent i could not be better off as the true dominant share it receives with manipulation would be no more than it received without manipulation. To justify the third transition, note that agent i must be allocated some resources at step k with manipulation. If $k = i$, this is trivial, and if $k > i$, this follows because otherwise k would not be the first step when agent i is strictly better off as we would have $u_i(\hat{\mathbf{A}}_i^{k-1}) = u_i(\hat{\mathbf{A}}_i^k) > u_i(\mathbf{A}_i^k) \geq u_i(\mathbf{A}_i^{k-1})$, where $\hat{\mathbf{A}}_i^k$ denotes the allocation to agent i at step k with manipulation. Thus, $\hat{x}_i^k > \hat{x}_i^{k-1}$, and the third transition now follows from Lemma 3. The last transition holds because \hat{x}_j^k satisfies the first constraint of the LP of step k . Thus, we conclude that $\hat{x}_j^k \geq x_j^k$ for all agents $j \leq k$. \square (Proof of Lemma 6)

Now, the mechanism satisfies DPO and thus allocates at least a k/n fraction of at least one resource at step k without manipulation. Let r be such a resource. Then the fraction of resource r allocated at step k with manipulation is

$$\hat{x}_i^k \cdot d_{ir} + \sum_{\substack{j \leq k \\ s.t. j \neq i}} \hat{x}_j^k \cdot d_{jr} > x_i^k \cdot d_{ir} + \sum_{\substack{j \leq k \\ s.t. j \neq i}} x_j^k \cdot d_{jr} = k/n.$$

To justify the inequality, note that $\hat{x}_i^k \cdot d_{ir} > x_i^k \cdot d_{ir}$ by Equation (1) (as agent i is strictly better off), and in addition $\hat{x}_j^k \geq x_j^k$ for every $j \leq k$. However, this shows that more than a k/n fraction of resource r must be allocated at step k with manipulation, which is impossible due to the third constraint of the LP. Hence, a successful manipulation is impossible, that is, DYNAMIC DRF is SP.

Finally, note that the LP has a linear number of variables and constraints, therefore the mechanism can be implemented in polynomial time. \square (Proof of Theorem 2)

5. RELAXING DYNAMIC PARETO OPTIMALITY

We saw (Theorem 1) that satisfying EF and DPO is impossible. We then explored an intuitive relaxation of EF. Despite the positive result (Theorem 2), the idea of achieving absolute fairness—as conceptualized by EF—in our dynamic setting is compelling.

As a straw man, consider waiting for all the agents to arrive and then using any EF static allocation mechanism. However, this scheme is highly inefficient, e.g., it is easy to see that one can always allocate each agent at least a $1/n$ share of its dominant resource (and other resources in proportion) as soon as it arrives and still maintain EF at every step. How much more can be allocated at each step? We put forward a general answer to this question using a relaxed notion of DPO that requires a mechanism to allocate as many resources as possible while ensuring that EF can be achieved in the future, but first we require the following definition. Given a step $k \in \{1, \dots, n\}$, define an allocation \mathbf{A} over the k present agents with demands $\mathbf{d}_{\leq k}$ to be *EF-extensible* if it can be extended to an EF allocation over all n agents with demands $\mathbf{d} = (\mathbf{d}_{\leq k}, \mathbf{d}_{>k})$, for all possible future demands $\mathbf{d}_{>k} \in \mathcal{D}^{n-k}$.

4'. *Cautious Dynamic Pareto optimality (CDPO)*. A dynamic allocation mechanism satisfies CDPO if at every step k , the allocation \mathbf{A}^k returned by the mechanism is not Pareto dominated by any other allocation \mathbf{A}' over the same k agents that is EF-extensible.

In other words, a mechanism satisfies CDPO if at every step it selects an allocation that is at least as generous as any allocation that can ultimately guarantee EF, irrespective of future demands.

At first glance, it may not be obvious that CDPO is indeed a relaxation of DPO. However, note that DPO requires a mechanism to allocate at least a k/n fraction of at least one resource r^* in the allocation \mathbf{A}^k at any step k , and thus to allocate at least a $1/n$ fraction of r^* to some agent i . Any alternative allocation that Pareto dominates \mathbf{A}^k must also allocate at least a $1/n$ fraction of r^* to

agent i . Hence, in order to ensure an EF extension over all n agents when all the future demands are identical to the demand of agent i , the alternative allocation must allocate at most a k/n fraction of r^* , as each future agent may also require at least a $1/n$ fraction of r^* to avoid envying agent i . It follows that the alternative allocation cannot Pareto dominate \mathbf{A}^k . Thus, the mechanism satisfies CDPO.

Recall that DYNAMIC DRF extends the water filling idea of the static DRF mechanism [8] to our dynamic setting. DYNAMIC DRF is unable to satisfy the original EF, because—to satisfy DPO—at every step k it needs to allocate resources until a k/n fraction of some resource is allocated. We wish to modify DYNAMIC DRF to focus only on competing with EF-extensible allocations, in a way that achieves CDPO and EF (as well as other properties).

The main technical challenge is checking when an allocation over k agents violates EF-extensibility. Indeed, there are uncountably many possibilities for the future demands $\mathbf{d}_{>k}$ over which an EF extension needs to be guaranteed by an EF-extensible allocation! Of course, checking all the possibilities explicitly is not feasible. Ideally, we would like to check only a small number of possibilities. The following lemma establishes that it is sufficient to verify that an EF extension exists under the assumption that all future agents will have the same demand vector that is moreover identical to the demand vector of one of the present agents.

LEMMA 7. *Let k be the number of present agents, $\mathbf{d}_{\leq k}$ be the demands reported by the present agents, and \mathbf{A} be an EF allocation over the k present agents. Then \mathbf{A} is EF-extensible if and only if there exists an EF extension of \mathbf{A} over all n agents with demands $\mathbf{d} = (\mathbf{d}_{\leq k}, \mathbf{d}_{>k})$ for all future demands $\mathbf{d}_{>k} \in \mathcal{D}'$, where $\mathcal{D}' = \{\langle \mathbf{d}_1 \rangle^{n-k}, \langle \mathbf{d}_2 \rangle^{n-k}, \dots, \langle \mathbf{d}_k \rangle^{n-k}\}$.*

To prove this lemma, we first introduce the notion of the *minimum EF extension*. Intuitively, the minimum EF extension is the “smallest” EF extension (allocating the least resources) of a given EF allocation to a larger set of agents. Formally, let \mathbf{A} be an EF allocation over a set of agents $S \subseteq N$ and \mathbf{A}^* be an EF extension of \mathbf{A} to a set of agents $T \subseteq N$ ($S \subseteq T$). Then \mathbf{A}^* is called the minimum EF extension of \mathbf{A} to T if for any EF extension \mathbf{A}' of \mathbf{A} to T , we have that \mathbf{A}' is an extension of \mathbf{A}^* . We show that the minimum EF extension exists and exhibits a simple structure.

LEMMA 8. *Let \mathbf{A} be an EF allocation over a set of agents $S \subseteq N$ and let x_i be the dominant share of agent $i \in S$ in \mathbf{A} . Let T be such that $S \subseteq T \subseteq N$ and let \mathbf{A}^* be an allocation over T with x_i^* as the dominant share of agent $i \in T$. Let $x_i^* = x_i$ for all $i \in S$, and $x_i^* = \max_{j \in S} y_j^j$ for all $i \in T \setminus S$, where $y_i^j = x_j \cdot \min_{r \in R} d_{jr} / d_{ir}$. Then \mathbf{A}^* is a minimum EF extension of \mathbf{A} to T .*

PROOF. For agent i with dominant share x_i to avoid envying agent j with dominant share x_j , there must exist $r \in R$ such that $x_i \cdot d_{ir} \geq x_j \cdot d_{jr}$, that is, $x_i \geq x_j \cdot d_{jr} / d_{ir}$. It follows that $x_i \geq x_j \cdot \min_{r \in R} d_{jr} / d_{ir}$, and thus the minimum dominant share is given by $y_i^j = x_j \cdot \min_{r \in R} d_{jr} / d_{ir}$. Now it is easy to argue that any EF extension \mathbf{A}' of \mathbf{A} over T must allocate at least an x_i^* dominant share to any agent $i \in T$, for both $i \in S$ and $i \in T \setminus S$, and thus \mathbf{A}' must be an extension of \mathbf{A}^* .

It remains to prove that \mathbf{A}^* is EF. First we prove an intuitive result regarding the minimum dominant share agent i needs to avoid envying agent j , namely y_i^j . We claim that for every $r \in R$,

$$y_i^j \cdot d_{ir} \leq x_j \cdot d_{jr}. \quad (2)$$

Indeed, to prevent agent i from envying agent j , we need to allocate no more than an $x_j \cdot d_{jr}$ fraction of resource r to agent i for

any $r \in R$. Formally, for any $r \in R$,

$$y_i^j \cdot d_{ir} = x_j \cdot \min_{r' \in R} \frac{d_{jr'}}{d_{ir'}} \cdot d_{ir} \leq x_j \cdot \frac{d_{jr}}{d_{ir}} \cdot d_{ir} = x_j \cdot d_{jr}.$$

Next we show that \mathbf{A}^* is EF, i.e., no agent i envies any agent j in \mathbf{A}^* . We consider four cases.

Case 1: $i \in S$ and $j \in S$. This case is trivial as \mathbf{A}^* is identical to \mathbf{A} over S and \mathbf{A} is EF.

Case 2: $i \in T \setminus S$ and $j \in S$. This case is also trivial because i receives at least a y_i^j fraction of its dominant resource.

Case 3: $i \in S$ and $j \in T \setminus S$. We must have $x_j = y_j^t$ for some $t \in S$. Agent i does not envy agent t in \mathbf{A} , and hence in \mathbf{A}^* . Thus, there exists a resource $r \in R$ such that $A_{ir}^* \geq A_{tr}^* \geq A_{jr}^*$, where the last step follows from Equation (2). Thus, agent i does not envy agent j .

Case 4: $i \in T \setminus S$ and $j \in T \setminus S$. Similarly to Case 3, let $x_j = y_j^t$ for some $t \in S$. Now $x_i \geq y_i^t$, so agent i does not envy agent t in \mathbf{A}^* . Thus, there exists a resource r such that $A_{ir}^* \geq A_{tr}^* \geq A_{jr}^*$, where again the last step follows from Equation (2).

Therefore, \mathbf{A}^* is an EF extension of \mathbf{A} over T and we have already established that any EF extension of \mathbf{A} over T must be an extension of \mathbf{A}^* . We conclude that \mathbf{A}^* is a minimum EF extension of \mathbf{A} over T . \square (Proof of Lemma 7)

It is not hard to see from the construction of the minimum EF extension that it not only exists, it is unique. We now prove Lemma 7.

PROOF OF LEMMA 7. The “only if” direction of the proof is trivial. To prove the “if” part, we prove its contrapositive. Assume that there exist future demand vectors $\hat{\mathbf{d}}_{>k} \in \mathcal{D}^{n-k}$ such that there does not exist any EF extension of \mathbf{A} to N with demands $\hat{\mathbf{d}} = (\mathbf{d}_{\leq k}, \hat{\mathbf{d}}_{>k})$. We want to show that there exists $\mathbf{d}'_{>k} \in \mathcal{D}'$ for which there is no EF extension as well.

Let $K = \{1, \dots, k\}$ and $N \setminus K = \{k+1, \dots, n\}$. Denote the minimum EF extension of \mathbf{A} to N with demands $\hat{\mathbf{d}}$ by \mathbf{A}^* . Let the dominant share of agent $i \in K$ in \mathbf{A} be x_i and the dominant share of agent $j \in N$ in \mathbf{A}^* be x_j^* .

No EF extension of \mathbf{A} over N with demands $\hat{\mathbf{d}}$ is feasible, hence \mathbf{A}^* must be infeasible too. Therefore, there exists a resource r such that $\sum_{i=1}^n x_i^* \cdot d_{ir} > 1$. Note that for every agent $j \in N \setminus K$, there exists an agent $i \in K$ such that $x_j^* = x_i \cdot \min_{r' \in R} d_{ir'} / d_{jr'}$, and hence $x_j^* \cdot d_{jr} \leq x_i \cdot d_{ir}$ by Equation (2). Taking the maximum over $i \in K$, we get that $x_j^* \cdot d_{jr} \leq \max_{i \in K} (x_i \cdot d_{ir})$ for every agent $j \in N \setminus K$. Taking $t \in \arg \max_{i \in K} (x_i \cdot d_{ir})$,

$$\begin{aligned} 1 &< \sum_{i=1}^n x_i^* \cdot d_{ir} = \sum_{i=1}^k x_i^* \cdot d_{ir} + \sum_{i=k+1}^n x_i^* \cdot d_{ir} \\ &\leq \sum_{i=1}^k x_i \cdot d_{ir} + (n-k) \cdot x_t \cdot d_{tr}. \end{aligned}$$

Consider the case where $\mathbf{d}'_{>k} = \langle \mathbf{d}_t \rangle^{n-k} \in \mathcal{D}'$. The minimum EF extension \mathbf{A}' of \mathbf{A} to N with demands $\mathbf{d}' = \langle \mathbf{d}_{\leq k}, \mathbf{d}'_{>k} \rangle$ allocates an x_i dominant share to every $i \in K$ (same as \mathbf{A}) and allocates exactly an x_t dominant share to every $j \in N \setminus K$. Thus, the fraction of resource r allocated in \mathbf{A}' is $\sum_{i=1}^k x_i \cdot d_{ir} + (n-k) \cdot x_t \cdot d_{tr} > 1$, implying that the minimum EF extension of $\mathbf{d}'_{>k}$ is infeasible. We conclude that there is no feasible EF extension for $\mathbf{d}'_{>k}$, as required. \square (Proof of Lemma 7)

The equivalent condition of Lemma 7 provides us with $k \cdot m$ linear constraints that can be checked to determine whether an allocation over k agents is EF-extensible. Using this machinery, we can write down a “small” linear program (LP) that begins with the

ALGORITHM 2: CAUTIOUS LP

Data: Demands d **Result:** Allocation \mathbf{A}^k at each step k $k \leftarrow 1;$ **while** $k \leq n$ **do**

$$\{x_i^k\}_{i=1}^k \leftarrow \text{Solution of the LP in the box below;}$$

$$A_{ir}^k \leftarrow x_i^k \cdot d_{ir}, \forall i \leq k;$$

$$k \leftarrow k + 1$$
end

Maximize M^k
 subject to
 $x_i^k \geq M^k, \forall i \leq k$
 $x_i^k \geq x_i^{k-1}, \forall i \leq k-1$
 $x_k^k \geq \max_{i \leq k-1} (x_i^{k-1} \cdot \min_{r \in R} d_{ir} / d_{kr})$
 $\sum_{i=1}^k x_i^k \cdot d_{ir} + (n-k) \cdot x_i^k \cdot d_{tr} \leq 1, \forall t \leq k, r \in R$

allocation chosen in the previous step (recall that the allocations are irrevocable), gives agent k a jump start so that it does not envy agents 1 through $k-1$, and then uses water filling to allocate resources similarly to DYNAMIC DRF, but subject to the constraint that the allocation stays EF-extensible. This intuition is formalized via the mechanism CAUTIOUS LP, which is given as Algorithm 2.

The mechanism’s third LP constraint jump-starts agent k to a level where it does not envy earlier agents, and the fourth LP constraint is derived from Lemma 7. To see why the mechanism satisfies CDPO, observe that if at any step k there is an EF-extensible allocation \mathbf{A}' that Pareto dominates the allocation \mathbf{A}^k returned by the mechanism, then (by Lemma 7) \mathbf{A}' must also satisfy the LP at step k . However, it can be shown that no allocation from the feasible region of the LP can Pareto dominate \mathbf{A}^k , otherwise we could redistribute some of the resources of the agent that is strictly better off to obtain a feasible allocation with a value of M^k that is higher than the optimal solution. It is also easy to see why intuitively CAUTIOUS LP is EF: the initial allocation to agent k achieves an EF allocation over the k agents, and water filling preserves EF because it always allocates to agents with minimum dominant share. It is equally straightforward to show that CAUTIOUS LP also satisfies SI. Establishing SP requires some work, but the proof is mainly a modification of the proof of Theorem 2 and appears in the full version.¹ We are therefore able to establish the following theorem, which formalizes the guarantees given by CAUTIOUS LP.

THEOREM 9. CAUTIOUS LP *satisfies SI, EF, CDPO, and SP, and can be implemented in polynomial time.*

6. EXPERIMENTAL RESULTS

We presented two potentially useful mechanisms, DYNAMIC DRF and CAUTIOUS LP, each with its own theoretical guarantees. Our next goal is to analyze the performance of both mechanisms on *real data*, for two natural objectives: the sum of dominant shares (the maxsum objective) and the minimum dominant share (the maxmin objective) of the agents present in the system.²

For the maxsum objective, we set the lower bound to be k/n at step k , and for the maxmin objective, we set the lower bound to be $1/n$ at any step. Note that these are provable lower bounds since

¹<http://www.cs.cmu.edu/~arielp/papers.html>

²Under a cardinal notion of utility where the dominant share of an agent is its utility, the sum of dominant shares (resp. the minimum dominant share) is the utilitarian (resp. egalitarian) social welfare.

both mechanisms satisfy SI. For upper bounds, we consider omniscient (hence unrealistic) mechanisms that maximize the objectives in an offline setting where the mechanisms have complete knowledge of future demands and only need to guarantee an EF extension on the real future demands rather than on all possible future demands. The comparison of CAUTIOUS LP with these offline mechanisms demonstrates the loss CAUTIOUS LP (an online mechanism) suffers due to the absence of information regarding the future demands, that is, due to its cautiousness. Because DYNAMIC DRF is not required to have an EF extension, the offline mechanisms are not theoretical upper bounds for DYNAMIC DRF, but our experiments show that they provide upper bounds in practice.

As our data we use traces of real workloads on a Google compute cell, from a 7 hour period in 2011 [1]. The workload consists of tasks, where each task ran on a single machine, and consumed memory and one or more cores; the demands fit our model with two resources. For various values of n , we sampled n random positive demand vectors from the traces and analyzed the value of the two objective functions under DYNAMIC DRF and CAUTIOUS LP along with the corresponding lower and upper bounds. We averaged over 1000 such simulations to obtain data points.

Figures 2(a) and 2(b) show the maxsum values achieved by the different mechanisms, for 20 agents and 100 agents respectively. The performance of our two mechanisms is nearly identical.

Figures 2(c) and 2(d) show the maxmin values achieved for 20 agents and 100 agents, respectively. Observe that DYNAMIC DRF performs better than CAUTIOUS LP for lower values of k , but performs worse for higher values of k . Intuitively, DYNAMIC DRF allocates more resources in early stages to satisfy DPO while CAUTIOUS LP cautiously waits. This results in the superior performance of DYNAMIC DRF in initial steps but it has fewer resources available and thus lesser flexibility for optimization in later steps, resulting in inferior performance near the end. In contrast, CAUTIOUS LP is later able to make up for its loss in early steps. Encouragingly, by the last step CAUTIOUS LP achieves near optimal maxmin value. For the same reason, unlike DYNAMIC DRF the maxmin objective value for CAUTIOUS LP monotonically increases as k increases in our experiments (although it is easy to show that this is not always the case).

7. DISCUSSION

We have presented a new model for resource allocation with multiple resources in dynamic environments that, we believe, can spark the study of dynamic fair division more generally. The model is directly applicable to data centers, clusters, and cloud computing, where the allocation of multiple resources is a key issue, and it significantly extends the previously studied static models. That said, the model also gives rise to technical challenges that need to be tackled to capture more realistic settings.

First, our model assumes positive demands, that is, each agent requires every resource. To see how the positive demands assumption plays a role, recall that achieving EF and DPO is impossible. We established that dropping DPO leads to the trivial mechanism EQUAL SPLIT, which satisfies the remaining three properties; this is also true for possibly zero demands. When we dropped EF, we observed that the trivial mechanism DYNAMIC DICTATORSHIP satisfies SI, DPO and SP, and we subsequently suggested the improved mechanism DYNAMIC DRF that satisfies DEF in addition to SI, DPO and SP. Surprisingly though, it can be shown that neither DYNAMIC DICTATORSHIP nor DYNAMIC DRF are SP under possibly zero demands.³ In fact, despite significant effort, we were

³In order to satisfy DPO under possibly zero demands,

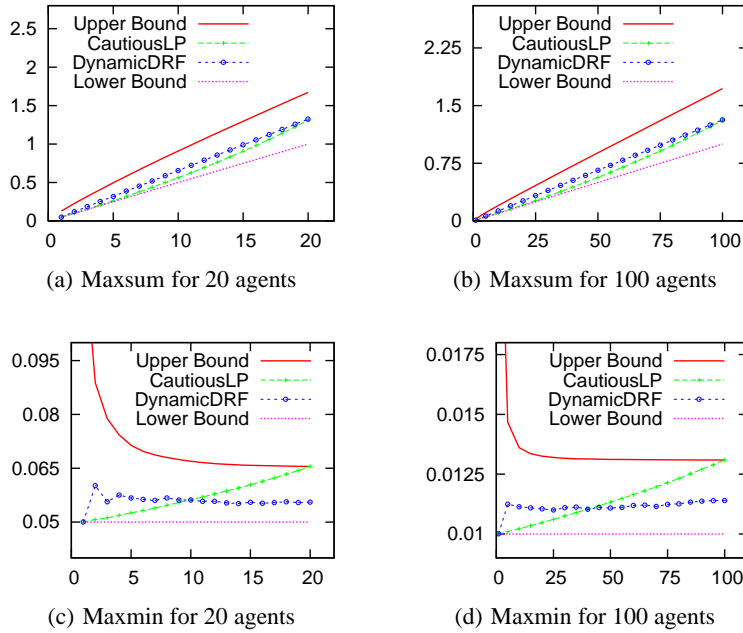


Figure 2: The maxsum and maxmin objectives as a function of the time step k , for $n = 20$ and $n = 100$.

unable to settle the question of the existence of a mechanism that satisfies SI, DPO and SP under possibly zero demands.

Second, our analysis is restricted to the setting of divisible tasks, where agents value fractional quantities of their tasks. Parkes et al. [11] consider the indivisible tasks setting, where only integral quantities of an agent’s task are executed, albeit in a static environment. It can be shown that even forward EF—the weakest of all EF relaxations considered in this paper—is impossible to achieve along with DPO under indivisible tasks. It remains open to determine which relaxations of EF are feasible in dynamic resource allocation settings with indivisible tasks.

Third, while our model of fair division extends the classical model by introducing dynamics, and our results can directly inform the design of practical mechanisms, we do make the assumption that agents arrive over time but do not depart which justifies our assumption of irrevocable allocations. In reality, agents may arrive and depart multiple times, and their preferences may also change over time (note that changing preferences can be modeled as a departure and simultaneous re-arrival with a different demand vector). Departures without re-arrivals are easy to handle; one can allocate the resources that become free in a similar way to allocations of entitlements, e.g., using water filling. However, departures with re-arrivals immediately lead to daunting impossibilities. Note though that mechanisms that were designed for static settings performed well in realistic (fully dynamic) environments [8], and it is quite likely that our mechanisms—which do provide theoretical guarantees for restricted dynamic settings—would yield even better performance in reality.

8. REFERENCES

[1] http://code.google.com/p/googleclusterdata/wiki/ClusterData2011_1.

DYNAMIC DICTATORSHIP and DYNAMIC DRF must be modified to continue allocating even when some resources become saturated.

[2] Y. Chen, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Truth, justice, and cake cutting. In *Proc. of 24th AAAI*, pages 756–761, 2010.

[3] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[4] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Reaching envy-free states in distributed negotiation settings. In *Proc. of 20th IJCAI*, pages 1239–1244, 2007.

[5] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. In *Proc. of SIGCOMM*, pages 1–12, 1989.

[6] D. Dolev, D. G. Feitelson, J. Y. Halpern, R. Kupferman, and N. Linial. No justified complaints: On fair sharing of multiple resources. In *Proc. of 3rd ITCS*, 2012. Forthcoming.

[7] A. Ghodsi, V. Sekar, M. Zaharia, and I. Stoica. Multi-resource fair queueing for packet processing. In *Proc. of SIGCOMM*, pages 1–12, 2012.

[8] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proc. of 8th NSDI*, pages 24–37, 2011.

[9] A. Gutman and N. Nisan. Fair allocation without trade. In *Proc. of 11th AAMAS*, pages 719–728, 2012.

[10] H. Moulin and R. Stong. Fair queueing and other probabilistic allocation methods. *Mathematics of Operations Research*, 27(1):1–30, 2002.

[11] D. C. Parkes, A. D. Procaccia, and N. Shah. Beyond Dominant Resource Fairness: Extensions, limitations, and indivisibilities. In *Proc. of 13th EC*, pages 808–825, 2012.

[12] A. D. Procaccia. Thou shalt covet thy neighbor’s cake. In *Proc. of 21st IJCAI*, pages 239–244, 2009.

[13] T. Walsh. Online cake cutting. In *Proc. of 3rd ADT*, pages 292–305, 2011.